

REPORT DOCUMENTATION PAGE

Form Approved
OMB NO. 0704-0188

Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimates or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE September 1996		3. REPORT TYPE AND DATES COVERED Final Report	
4. TITLE AND SUBTITLE Virtual Humans for Animation, Ergonomics, and Simulation				5. FUNDING NUMBERS DAAH04-94-G-0402	
6. AUTHOR(S) Dr. Norman I. Badler					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Pennsylvania Computer & Information Science Department Center for Human Modeling & Simulation Phila., PA 19104-6389				8. PERFORMING ORGANIZATION REPORT NUMBER DAAH04-94-G-0402	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U. S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211				10. SPONSORING / MONITORING AGENCY REPORT NUMBER ARO 32931.1-MA	
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
12 a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.				12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This project describes state of the art computation speed and control methods needed to portray #D virtual humans suitable for real interactive applications. Then focus is given to the particular approach taken at the University of Pennsylvania with the Jack system. Various aspects of real-time virtual humans are considered, such as appearance and motion, interactive control, autonomous action, gesture, attention, locomotion, and multiple individuals. The underlying architecture consists of a sense-control-act structure that permits reactive behaviors to be locally adaptive to the environment, and a "PaT-Net" parallel finite-state machine controller that can be used to drive virtual humans through complex tasks.					
14. SUBJECT TERMS Computer Graphics				15. NUMBER OF PAGES 15	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL		

Virtual Humans for Animation, Ergonomics, and Simulation
FINAL REPORT
~~DAAH04-94-G-0402 (DMSO)~~
September 1996

Norman Badler
Center for Human Modeling and Simulation
University of Pennsylvania
Philadelphia, PA 19104-6389
215-898-5862
badler@central.cis.upenn.edu

Abstract

The last few years have seen great maturation in the computation speed and control methods needed to portray 3D virtual humans suitable for real interactive applications. We first describe the state of the art, then focus on the particular approach taken at the University of Pennsylvania with the *Jack* system. Various aspects of real-time virtual humans are considered, such as appearance and motion, interactive control, autonomous action, gesture, attention, locomotion, and multiple individuals. The underlying architecture consists of a sense-control-act structure that permits reactive behaviors to be locally adaptive to the environment, and a "PaT-Net" parallel finite-state machine controller that can be used to drive virtual humans through complex tasks.

1 Virtual Humans

Only fifty years ago, computers were barely able to compute useful mathematical functions. Twenty-five years ago, enthusiastic computer researchers were predicting that all sorts of human tasks from game-playing to automatic robots that travel and communicate with us would be in our future. Today's truth lies somewhere in-between. We have balanced our expectations of complete machine autonomy with a more rational view that machines should assist people to accomplish meaningful, difficult, and often enormously complex tasks. When those tasks involve human interaction with

20001122 143

the physical world, computational representations of the human body can be used to escape the constraints of presence, safety, and even physicality.

Virtual humans are computer models of people that can be used

- as substitutes for “the real thing” in *ergonomic* evaluations of computer-based designs for vehicles, work areas, machine tools, assembly lines, etc., *prior to the actual construction of those spaces*;
- for *embedding real-time representations of ourselves or other live participants* into virtual environments.

Recent improvements in computation speed and control methods have allowed the portrayal of 3D humans suitable for interactive and real-time applications. There are many reasons to design specialized human models that individually optimize character, performance, intelligence, and so on. Many research and development efforts concentrate on one or two of these criteria.

In the efforts that we describe here, we cross several domains which in turn build from various interrelated facets of human beings (Fig. 1):

- **Human Factors Analysis:** Human size, capabilities, behavior, and performance affects work in and use of designed environments.
- **Real-Time Agents and Avatars:** People come from different cultures and have different personalities; this richness and diversity must be reflected in virtual humans since it influences appearance as well as reaction and choice.
- **Instruction Understanding and Generation:** Humans communicate with one another within a rich context of shared language, senses, and experience and this needs to be extended to computer-generated agents and avatars.
- **Bio-Medical Simulation:** The human machine is a complex of physical structures and functions; to understand human behavior, physiological responses, and injuries we need to represent biological systems.
- **Motion and Shape Analysis:** Understanding what we perceive when we see or sense the world leads to models of the physical world (physics) and the geometric shapes and deformations of objects.

From these virtual humans research areas, many current, emergent, or future major applications are enabled:

- **Engineering:** Analysis and simulation for virtual prototyping and simulation-based design.
- **Virtual-Conferencing:** Efficient tele-conferencing using virtual representations of participants to reduce transmission bandwidth requirements.
- **Interaction:** Agents and avatars that insert real-time humans into virtual worlds with virtual reality.

- Monitoring: Acquiring, interpreting, and understanding shape and motion data on human movement, performance, activities, or intent.
- Virtual Environments: Living and working in a virtual place for visualization, analysis, training, or just the experience.
- Games: Real-time characters with actions and personality for fun and profit.
- Training: Skill development, team coordination, and decision-making.
- Education: Distance mentoring, interactive assistance, and personalized instruction.
- Military: Battlefield simulation with individual participants, team training, and peace-keeping operations.
- Design/Maintenance: Design for access, ease of repair, safety, tool clearance, visibility, and hazard avoidance.

Besides general industry-driven improvements in the underlying computer and graphical display technologies themselves, virtual humans will enable quantum leaps in applications requiring personal and live participation.

In building models of virtual humans, there are varying notions of *virtual fidelity*. Understandably, these are application dependent. For example, fidelity to human size, capabilities, and joint and strength limits are essential to some applications such as design evaluation; whereas in games, training, and military simulations, temporal fidelity (real-time behavior) is essential. In our efforts we have attacked both.

Understanding that different applications require different sorts of virtual fidelity leads to the question of what makes a virtual human "right"?

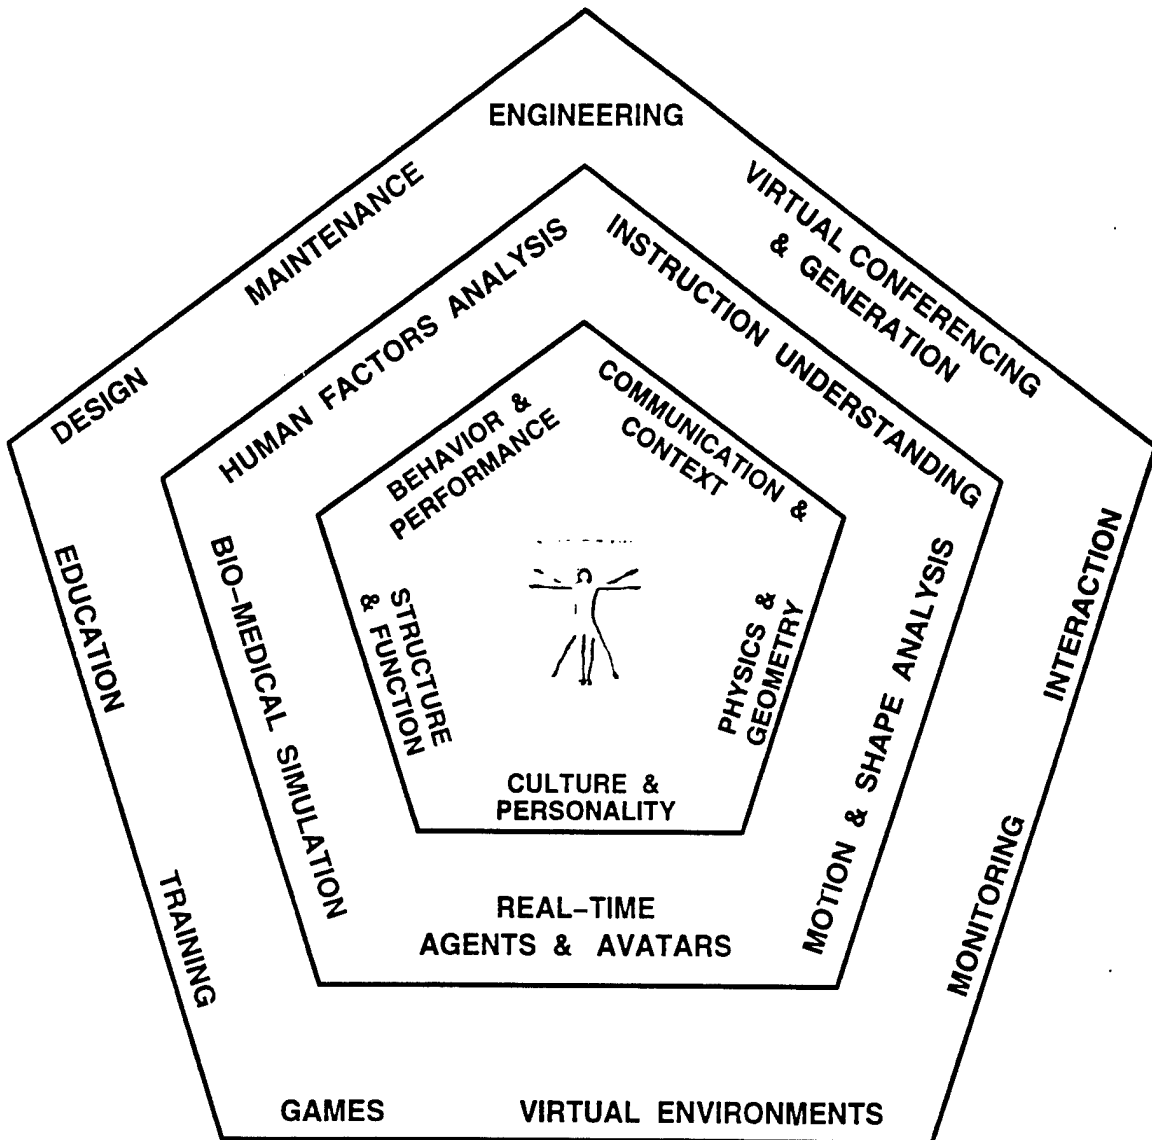
- What do you want to do with it?
- What do you want it to look like?
- What characteristics are important to success of the application?

Unfortunately the state of research in virtual humans is not as advanced as to make the proper selection a matter of buying off-the-shelf systems. There are gradations of fidelity in the models: some models are very advanced in a narrow area but lack other desirable features.

In a very general way, we can characterize the state of virtual human modeling along at least five dimensions:

- Appearance: Cartoon shape - - - - - + - - > Physiologically accurate model
- Function: Cartoon actions - - - - - + - > Human limitations
- Time: Off-line generation - - - - - + - > Real-time production
- Autonomy: Direct animation - - - - - + - - > Intelligent
- Individuality: Specific person - - - + - - - - > Varying personalities

Virtual Humans



The arrows and hash marks are meant to be qualitative indicators of where we think usable technology exists today. Understanding that the arrows can actually extend an undetermined distance to the right, the idea is nonetheless being conveyed that we (and others) have proceeded rather far beyond the individual rendering of still frames as realized by traditional hand animation or even computer assisted cartoon animation. If we need to invoke them, the *appearance* of increasingly accurate physiologically- and biomechically-grounded human models may be obtained. We can create virtual humans with *functional* limitations that go beyond cartoons into instantiations of known human factors data. Animated virtual humans can be created in human *time* scales through motion capture or computer synthesis. Virtual humans are also beginning to exhibit the early stages of *autonomy* and intelligence as they react and make decisions in novel, changing environments rather than being forced into fixed movements. Finally, rather preliminary investigations are underway to create characters with *individuality* and personality who react to and interact with other real or virtual people [Bate92a, Bate92b, Cass94, Maes95, Perl96, Rous96].

The University of Pennsylvania has been very actively engaged in research and development of human-like simulated figures. Our interest in human simulation is not unique, but the complex of activities surrounding our approach is. The framework for our research is a software system called *Jack* [Badl93b]. *Jack* is an interactive system for definition, manipulation, animation, and performance analysis of virtual human figures. Our philosophy has led to a particular realization of a virtual human model that pushes the above five dimensions toward the right:

- can be substituted for live individuals for workspace or cockpit evaluation.
- demonstrates various (useful) human limitations, constraints, and capabilities.
- may be moved live (in real-time) by position and orientation information or other motion generators such as *walk-to* or *look-at*.
- may have its actions synthesized by a program so that it can make its own decisions, navigate spaces, and so on.
- represents "anyone" rather than a single specific person or character.

Virtual humans are different than simplified cartoon and game characters. What are the characteristics of this difference and why are virtual humans more difficult to construct? After all, anyone who goes to the movies can see marvelous synthetic characters (aliens, toys, dinosaurs, etc.), but they have been created typically for one scene or one movie and are not meant to be re-used (except possibly by the animator – and certainly not by the viewer). The difference lies in the *interactivity* and *autonomy* of virtual humans. What makes a virtual human *human* is not just a well-executed exterior design but movements, reactions, and decision-making which appear "natural," appropriate, and contextually-sensitive.

2 Agents and Avatars

We will consider an *agent* to be a virtual human figure representation that is created and controlled by computer programs. An *avatar* is a virtual human controlled by a live participant. The principal issues roughly follow the dimensions cited above: *appearance and motion*, mechanisms of *control for interactivity and autonomy*, including gesture, attention, and locomotion, and multi-agent *interaction, cooperation, and coordination*.

2.1 Appearance and Motion

Avatars can be portrayed visually as 2D icons, cartoons [Kurl96], composited video, 3D shapes, or full 3D bodies [Badl93a, Stan94, Robe96]. We are mostly interested in portraying human-like motions, so naturally tend toward the more realistic surface and articulation structures. In general, we prefer to design motions for highly articulated models and then reduce both the model detail and the articulatory detail as demanded by the application [Gran95].

Along the appearance dimension, the *Jack* figure has developed as a polygonal model with rigid segments and joint motions and limits accurate enough for ergonomics evaluations [Badl93b]. For real-time avatar purposes, simpler geometry can be used provided that the overall impression is one of a task-relevant figure. Thus a soldier model with 110 polygons is acceptable if drawn small enough and colored and/or texture mapped to be recognized as a soldier. On the other hand, a vehicle occupant model must show accurate and visually continuous joint geometry under typical motions. It must be both an acceptable occupant surrogate as well as a pleasing model for the non-technical viewer – who may be used to going to the movies to see the expensive special effects figures. Our “smooth body” [Azuo94] was developed using free-form deformation techniques [Sede86] to aid in the portrayal of visually appealing virtual humans.

The motions manifest in the avatar may arise from various sources:

- Motion capture from direct live video
- Motion capture from sensors
- Pre-stored motion data
 - as 2D sprites
 - as 3D global transformations
 - as 3D local (joint) transformations
- Motion synthesis
 - joint angle interpolation
 - inverse kinematics
 - dynamics
 - other generators (e.g. locomotion, faces)

In general, we will not consider 2D or purely video presentations of avatars, rather we will concentrate on avatars that more-or-less mimic human structure.

The distinction between “synthesized” motions and the other types is roughly that the former generate transformations for more than one joint at a time. Thus, for example, we store a time series of joint angle changes (per joint) in *channelsets* so that specific motions can be re-played under real-time constraints [Gran95]. No deviation from the pre-stored local transformations are allowed, although the whole body may be re-oriented or the playback speed varied. In a particularly effective modification of this technique, Perlin adds periodic noise to real-time joint transformations to achieve greater movement variability, animacy, and motion transitions [Perl95].

In a motion synthesizer, a small number of parameters control a much greater number of joints, for example:

- end effector position and orientation can control joints along an articulated chain [Zhao94, Koga94, Tola96],
- a path or footsteps can control leg and foot rotations through a locomotion model [Gira85, Ko96],
- a balance constraint can be superimposed on gross body motions [Badl93b, Ko96],
- dynamics calculations can move joints subject to arbitrary external and internal applied forces [Kokk96, Meta96],
- secondary motions can enhance a simpler form [Perl95, Hodg95].

The relative merits of pre-stored and synthesized motions must be considered when implementing virtual humans. The advantages to pre-stored motions are primarily speed of execution and algorithmic security (by minimizing computation). The major advantages to synthesis are the reduced parameter set size (and hence less information that needs to be acquired or communicated) and the concomitant generalized motion control: walk, reach, look-at, etc. The principal disadvantages to pre-stored motion are their lack of generality (since every joint must be controlled explicitly) and their lack of anthropometric extensibility (since changing joint-to-joint distances will change the computed locations of end effectors such as feet, making external constraints and contacts impossible to maintain). The disadvantages to synthesis are the difficulty of inventing natural-looking motions and the potential for positional disaster if the particular parameter set or code should have no solution, fail to converge on a solution, or just compute a poor result. In particular, we note that inverse kinematics is not in itself an adequate model of human *motion* – it is just a local positioning aid [Badl93b, Koga94]. The issue of building adequate human motion synthesis models is a wide open and complex research topic.

Since accurate human motion is difficult to synthesize, motion capture is a popular alternative, but one must recognize its limited adaptability and subject specificity. Although a complex motion may be used as performed, say in a CD-ROM game or as the source material for a (non-human) character animation, the motions may be best

utilized if segmented into motion “phrases” that can be named, stored, and executed separately, and possibly connected with each other via transitional (non-captured) motions [Brud95, Rose96]. Several projects have used this technique to interleave “correct” human movements into simulations that control the order of the choices. While 2D game characters have been animated this way for years – using pre-recorded or hand animated sequences for the source material – recently the methods have graduated to 3D whole body controls suitable for 3D game characters, real-time avatars, and military simulations that include individual synthetic soldiers [Prat94, Gran95, Bost96].

2.2 Control for Interactivity

Whichever motion generation technique is used, there must be a way of triggering the desired activity in the avatar. Specifying the motion can be as simple as direct sensor tracking (where each joint is driven by a corresponding sensor input), end effector tracking (where inverse kinematics or other behaviors generate the “missing” joint data), or external invocation via menu, speech, or button selection of the actions (whether then synthesized or interpreted from pre-stored data). The interesting observation is that the only mechanism available to an “unencumbered” participant is actually speech! Any other avatar control mechanism requires either a hands-on device (mouse, keyboard, glove input), or else external sensors and a limited field of movement. While there is considerable progress in using computer vision techniques to capture human motion [Azuo94, Essa95, DeCa96, Kaka96], both user mobility and movement generality are still in the future. Our intention is not to promote speech input *per se*, but to use this observation to promote (in Section 3 a *language-centered* view of action “triggering” augmented and elaborated by lower-level motion synthesis or playback. Although textual instructions can describe and trigger actions, details need not be explicated communicated. Thus the agent/avatar architecture must include semantic interpretation of instructions and even a lower reactive level within the movement generators that allows motion generality and environmental context-sensitivity.

2.3 Control for Autonomy

Providing a virtual human with human-like reactions and decision-making is more complicated than controlling its joint motions from captured or synthesized data. Here is where we engage the viewer with the character’s personality and demonstrate its skill and intelligence in negotiating its environment, situation, and other agents. This level of performance requires significant investment in decision-making tools. We presently use a two level architecture:

- to optimize reactivity to the environment at the lower level (for example, in the choice of footsteps for locomotion through the space) [Reic94, Ko96];
- to execute parametrized scripts or plan complex task sequences at the higher level (for example, choosing which room to search in order to locate an object or

another agent, or outlining the primary steps that must be followed to perform a particular task) [Moor95, Badl96].

The architecture is built on Parallel Transition Networks (PaT-Nets) [Badl93b]. Nodes represent executable processes, edges contain conditions which when true cause transitions to another node (process), and a combination of message passing and global memory provide coordination and synchronization across multiple parallel processes. Elsewhere we have shown how this architecture can be applied to the game of "Hide and Seek" [Badl96], to two person animated conversation [Cass94], or to simulated emergency medical care [Chi96]. Currently we are using this architecture to construct appropriate gestural responses from a synthetic agent, create appropriate visual attention during high-level task execution, manage locomotion tasks, and study multi-agent activity scheduling.

2.4 Gesture Control

Human arms serve (at least) two separate functions: they permit an agent/avatar to change the local environment through dextrous activities by reaching for and grasping (getting control over) objects [Gour89, Douv96], and they serve social interaction functions by augmenting the speech channel with communicative emblems, gestures and beats [Cass94].

For the first function, a consequence of human dexterity and experience is that we are rarely told how to approach and grasp an object. Rather than have our virtual humans learn – through direct experience and errors – how to grasp an object, we provide assistance through an object-specific relational table (OSR). Developed from ideas about *object-specific reasoning* [Levi96], the OSR has fields for each graspable site (in the *Jack* sense of an oriented coordinate triple) describing the appropriate handshape, grasp approach direction, and most importantly, its function or purpose. The OSR is manually created for graspable objects and allows an agent to look up an appropriate grasp site given a purpose, use the approach vector as guidance for the inverse kinematics directives that move the arm, and know which handshape is likely to result in reasonable finger placement. The hand itself is closed on the object through local geometry information and collision detection.

The second function of gestures is non-verbal communication. Thus gestures can be metaphors for actual objects, give indicators (via pointing) of location or participants in a virtual space around the speaker, or augment the speech signal with beats for added emphasis [Cass94]. Currently we are working on embedding culture-specific and even individual personality gesture variations. The potential interference between practical and gestural functions is leading to a resource-based priority model to resolve conflicts.

Given that arm control for avatars requires fast position and orientation of the hands for either reaching or gestural function, fast computation of arm joint angles is essential. In recent work we have pushed beyond iterative inverse kinematics [Zhao94] to analytic

formulas that can easily keep up with a live performance or a motion synthesizer outputting end effector position and orientation streams [Tola96]. By extending this idea to the whole body, multiple individuals (3-10 on an SGI RE2) may be controlled in real-time by arbitrary end-effector and global body data alone [Zhao96].

2.5 Attention Control

A particularly promising connection is underway to connect PaT-Nets into other high level "AI-like" planning tools for improved cognitive performance of virtual humans. As tasks are generated for the *Jack* figure, they are entered into a task queue. An *attention resource manager* [Chop95] scans this queue for current and future visual sensing requirements, and directs *Jack's* eye gaze (and hence head movement) accordingly. For example, if the agent is being told to "remove the power supply," parallel instructions are generated to locomote to the power supply area and attend to specific visual attention tasks such as searching for the power supply, scanning for potential moving objects, and periodically watching for obstacles near the feet. Note that normally none of this attentional information appears explicitly in the task-level instruction stream.

2.6 Locomotion with anticipation

In order to interact with a target object, an agent must determine that it is not within a suitable distance and must therefore locomote to a task-dependent position and orientation prior to the initiation of the reach and grasp. Such a decision is readily made by embedding it in a PaT-Net representing potential actions that enable the specified action. Moreover, the locomotion process itself uses the two level architecture: at the lowest level the agent or avatar gets a goal and an explicit list of objects to be avoided; the other level encapsulates locomotion states and decisions about transitions. For example, the agent could be walking, hiding, searching, or chasing. If walking, then transitions can be based on evaluating the best position of the foot relative to the goal and avoidances. If hiding, then assessments about line of sight between virtual humans are computed. If searching, then a pattern for exhaustively checking the local geometry is invoked. Finally, if chasing, then the goal is the target object: but if the target goes out of sight, the last observed position is used as an interim goal.

2.7 Multi-agent task allocation

By encapsulating virtual human activities in PaT-Nets, we can interactively control the assignment of tasks to agents. A menu or program binds actions to individuals, who then execute the PaT-Net processes. Since the processes have the power to query the environment and other agents before starting to execute, multi-agent synchronization and coordination can be modeled. Thus an agent can start a task when another signals that the situation is ready, or one agent can lead another in a shared task. The latter

would be especially useful when an avatar works with a simulated agent to perform a two-person task. One virtual human is designated as the "leader" (typically the avatar, so the live participant is in control) and the other the "follower." The follower's timing and motion are performed after each time-stepped motion of the leader. (The reverse situation, where the agent leads the avatar, may be needed for training and educational applications.) These are clearly the first steps toward a virtual social architecture.

We developed a prototype system for agent task assignment to evaluate a multi-function aircraft maintenance equipment cart ("MASS"). The user specifies tasks for an agent, and the agent accepts tasks for which it is both qualified and responsible. The tasks can be queued in advance, and are executed as prior tasks are completed or other agent or environment conditions obtain.

Once we can generate and control multiple agents and avatars, many social and community issues arise including authentication of identity, capabilities, permissions, social customs, transference of object control, sharing behaviors, coordinating group tasks, etc. Underlying technology to share interactive experience will depend on distributed system protocols and communication technology, client workstation performance, avatar graphics, and so on. Many of these issues are being addressed by other *ad hoc* groups, such as *Living Worlds*, *Open Community*, and *Universal Avatars*. Having two avatars "shake hands" is considered the first stage of a social encounter requiring significant attention to the details of avatar interaction, body representation, and action synchronization. Assuming that the communications can be done fast enough (a big assumption), our avatars should be able to reach for each other's hand, detect a collision/connection, and then allow the follower avatar to position his/her hand according to the leader's spatial position. Indeed, such a demonstration has already been readily constructed by Stansfield at Sandia National Labs with *Jack* avatars, in-house network communication software, head-mounted displays, and end effector position/orientation sensors on the participants.

3 Connecting Language and Animation

Even with a powerful set of motion generators, a challenge remains to provide effective and easily learned user interfaces to control, manipulate and animate virtual humans. Interactive point and click systems such as *Jack* work now, but with a cost in user learning and menu traversal. Such interfaces decouple the human participant's instructions and actions from the avatar through a narrow and *ad hoc* communication channel of hand and finger motions. A direct programming interface, while powerful, must be rejected as an off-line method that moreover requires specialized computer programming understanding and expertise. The option that remains is a language-based interface.

Perhaps not surprisingly, instructions for people are given in natural language augmented with graphical diagrams and occasionally, animations. Recipes, instruction manuals, and interpersonal conversations use language as the medium for conveying process and action. While our historic interest in instructions has been on creating

animations from instructions [Badl90, Badl93b, Webe95], we have recently begun to examine the inverse process, namely, generating text from the PaT-Net representations of animations. The purpose is primarily to help automate the production of aircraft maintenance instruction orders (manuals) in conjunction with the animation of the tasks themselves. The expectation is that the synthesized *text* material ought to reflect the proper execution of the tasks (which can be *visually* verified through the animation) and will have consistency across the entire document. By the same principles, being able to process the textual instructions will aid in discovering ambiguities, omitted steps, or inappropriate terminology.

The key to linking language and animation lies in constructing a semantic representation of actions, objects, and agents which is simultaneously suitable for execution (animation) as well as natural language expression. We have called this *implementable semantics*: the representation must have the power of a (parallel) programming language which drives a simulation (in a context of a given set of objects and agents), and yet supports the enormous range of expression, nuance, and manner offered by language. The details of this Parameterized Action Representation (PAR) – which involves PaT-Nets as an implementation language – are currently being developed.

4 Conclusions

The future holds great promise for the virtual humans who will populate our virtual worlds. They will provide economic benefits by helping designers early in the product design phases to produce more human-centered vehicles, equipment, assembly lines, manufacturing plants, and interactive systems. Virtual humans will enhance the presentation of information through training aids, virtual experiences, and even teaching and mentoring. And Virtual humans will help save lives by providing surrogates for medical training, surgical planning, and remote telemedicine. They will be our avatars on the Internet and will portray ourselves to others, perhaps as we are or perhaps as we wish to be. They may help turn cyberspace into a real, or rather virtual, community.

Acknowledgments

This research is partially supported by DARPA DAMD17-94-J-4486; U.S. Air Force DEPTH through Hughes Missile Systems F33615-91-C-0001; U.S. Air Force through BBN F33615-91-D-0009/0008; U.S. Air Force DAAH04-95-1-0151; DMSO DAAH04-94-G-0402; ONR through Univ. of Houston K-5-55043/3916-1552793; ARO DURIP DAAH04-95-1-0023; Army AASERT DAAH04-94-G-0220; DARPA AASERT DAAH04-94-G-0362; NSF IRI95-04372; National Library of Medicine N01LM-43551; National Institute of Standards and Technology 60 NANB6D0149; and JustSystem Japan.

References

- [Azu094] Azuola, F., Badler, N., Ho, P.-H., Kakadiaris, I., Metaxas, D. and Ting, B. Building anthropometry-based virtual human models. Proc. IMAGE VII Conference, Tucson, AZ (1994).
- [Badl93a] Badler, N., Hollick, M. and Granieri, J. Real-time control of a virtual human using minimal sensors. *Presence* 2(1), pp. 82-86 (1993).
- [Badl93b] Badler, N., Phillips, C. and Webber, B. *Simulating Humans: Computer Graphics Animation and Control*. Oxford University Press, New York, NY (1993).
- [Badl96] Badler, N., Webber, B., Becket, W., Geib, C., Moore, M., Pelachaud, C., Reich, B. and Stone, M. Planning for animation. To appear in *Computer Animation*, ed. N. Magnenat-Thalmann and D. Thalmann. Prentice-Hall (1996).
- [Badl90] Badler, N., Webber, B., Kalita, J. and Esakov, J. Animation from instructions. In *Making Them Move: Mechanics, Control, and Animation of Articulated Figures*. (N. Badler, B. Barsky, and D. Zeltzer, eds.) Morgan-Kaufmann (1990), pp. 51-93.
- [Bate92a] Bates J. The role of emotion in believable agents. *Communications of the ACM* 37(7), pp. 122-125 (1994).
- [Bate92b] Bates J., Loyall A., Reilly W. Integrating reactivity, goals, and emotion in a broad agent. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum, Hillsdale, NJ. pp. 696-701 (1992).
- [Bost96] Boston Dynamics, Inc., Cambridge, MA. BDI-Guy. Software system. 1996.
- [Brud95] Bruderlin, A. and Williams, L. Motion signal processing. *Computer Graphics. Annual Conference Series*, ACM (1995), 97-104.
- [Cass94] Cassell, J., Pelachaud, C., Badler, N., Steedman, M., Achorn, B., Becket, W., Douville, B., Prevost, S. and Stone, M. Animated Conversation: Rule-based generation of facial expression, gesture and spoken intonation for multiple conversational agents. *Computer Graphics, Annual Conference Series*, ACM (1994), 413-420.
- [Chi96] Chi, D., Webber, B., Clarke, J. and Badler, N. Casualty modeling for real-time medical training. *Presence* 5(4), pp. 359-366 (1996).
- [Chop95] Chopra, S. Strategies for simulating direction of gaze and attention. Internal report, Center for Human Modeling and Simulation, University of Pennsylvania (1995).
- [DeCa96] DeCarlo, D. and Metaxas, D. The integration of optical flow and deformable models with applications to human face shape and motion estimation. In *Proceedings CVPR '96*, pp. 231-238 (1996).
- [Douv96] Douville, B., Levison, L. and N. Badler, N. Task level object grasping for simulated agents. *Presence* 5(4) (1996) pp. 416-430.

- [Essa95] Essa, I. and Pentland, A. Facial expression recognition using a dynamic model and motion energy. In *Proceedings of the International Conference on Computer Vision*, Cambridge, MA (1995).
- [Gira85] Girard, M. and Maciejewski, A. Computational modeling for the computer animation of legged figures. *Computer Graphics* 19(3), pp. 263-270 (1985).
- [Gour89] Gourret, J.-P., Magnenat-Thalmann, N. and Thalmann, D. Simulation of object and human skin deformations in a grasping task. *ACM Computer Graphics*, 23(3) (1989) pp. 21-30.
- [Gran95] Granieri, J., Crabtree, J. and Badler, N. Off-Line production and real-time playback of human figure motion for 3D virtual environments. *VRAIS Conference*, IEEE Press, 1995.
- [Hodg95] Hodgins, J., Wooten, W., Brogan, D. and O'Brien, J. Animating human athletics. *Computer Graphics*, Annual Conference Series, ACM (1995), pp. 71-78.
- [Kaka96] Kakadiaris, I. and Metaxas, D. Model-based estimation of 3D human motion with occlusion based on active multi-viewpoint selection. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 81-87, San Francisco, CA (June 1996).
- [Ko96] Ko, H. and Badler, N. Animating human locomotion in real-time using inverse dynamics, balance and comfort control. *IEEE Computer Graphics and Applications* 16(2), pp. 50-59 (March 1996).
- [Koga94] Koga, Y., Kondo, K., Kuffner, J. and Latombe, J.-C. Planning motions with intentions. *Computer Graphics*, Annual Conference Series, ACM (1994), 395-408.
- [Kokk96] Kokkevis, E., Metaxas, D. and Badler, N. User-controlled physics-based animation for articulated figures. *Computer Animation. Conf. Proceedings* (1996).
- [Kurl96] Kurlander, D., Skelly, T., and Salesin, D. Comic chat. *Computer Graphics*, Annual Conference Series, ACM (1996), 225-236.
- [Levi96] Levison, L. Connecting planning and acting via object-specific reasoning. PhD thesis. University of Pennsylvania (1996).
- [Maes95] Maes P., Darrell T., Blumberg B., Pentland A. The ALIVE system: Full-body interaction with autonomous agents. In *Proceedings Computer Animation '95*, N. Magnenat-Thalmann and D. Thalmann D.(eds.), Geneva, Switzerland, IEEE Computer Society, Los Alamitos, CA, pp. 11-18 (1995).
- [Meta96] Metaxas, D. *Physics-Based Modeling: Applications to Computer Vision, Graphics and Medical Imaging*. Kluwer, Boston, MA (1996).
- [Moor95] Moore, M. B., Geib, C. W. and Reich, B. D. Planning and terrain reasoning. *AAAI Spring Symposium on Integrated Planning Applications*, Stanford CA (March 1995). (Also appears as Technical Report MS-CIS-94-63, University of Pennsylvania).

- [Perl95] Perlin, K. Real time responsive animation with personality. *IEEE Trans. on Visualization and Computer Graphics* 1(1), pp. 5-15 (1995).
- [Perl96] Perlin, K. and Goldberg, A. Improv: A system for scripting interactive actors in virtual worlds. *Computer Graphics, Annual Conference Series*, ACM (1996), 205-216.
- [Prat94] Pratt, D., Barham, P., Locke, J., Zyda, M., Eastman, B., Moore, T., Biggers, K., Douglass, R., Jacobsen, S., Hollick, M., Granieri, J., Ko, H. and Badler, N. Insertion of an articulated human into a networked virtual environment. *Proc. of the Conference on AI, Simulation and Planning in High Autonomy Systems*, University of Florida, Gainesville (1994).
- [Reic94] Reich, B., Ko, H., Becket, W. and Badler, N. Terrain reasoning for human locomotion. *Proceedings of Computer Animation '94*. Geneva. IEEE Computer Society Press (1994), 996-1005.
- [Robe96] Robertson, B. Best behaviors. Digital Magic, Supplement to *Computer Graphics World*, pp. S12-S19 (August 1996).
- [Rose96] Rose, C., Guenter, B., Bodenheimer, B. and Cohen, M. Efficient generation of motion transitions using spacetime constraints. *Computer Graphics, Annual Conference Series*, ACM (1996), 147-154.
- [Rous96] Rousseau, D. and Hayes-Roth, B. Personality in synthetic agents. Stanford Knowledge Systems Laboratory, Report KSL-96-21 (1996).
- [Sede86] Sederberg, T. and Parry, S. Free-form deformation of solid geometric models. *Computer Graphics* 20(4), pp. 151-160 (August 1986).
- [Stan94] Stansfield, S. Distributed virtual reality simulation system for situational training. *Presence: Teleoperators and Virtual Environments* 3(4), pp. 360-366 (1994).
- [Tola96] Tolani, D. and Badler, N. Real-time inverse kinematics for the human arm. *Presence* 5(4), pp. 393-401 (1996).
- [Webe95] Webber, B., Badler, N., Di Eugenio, B., Geib, C., Levison, L. and Moore, M. Instructions, intentions and expectations. *Artificial Intelligence Journal* 73, pp. 253-269 (1995).
- [Zhao94] Zhao, J. and Badler, N. Inverse kinematics positioning using nonlinear programming for highly articulated figures. *ACM Transactions on Graphics* 13(4) (1994) pp. 313-336.
- [Zhao96] Zhao, X. Kinematic control of human postures for task simulation. PhD thesis, University of Pennsylvania (1996).